

INVENTOR: George Q. CHEN

**METHOD AND APPARATUS FOR RECOVERING DEPTH USING
MULTI-PLANE STEREO AND SPATIAL PROPAGATION**

Cross-Reference To Related Application

This application generally relates to the teachings of U.S. Patent Application No. 09/825,266, entitled "Methods And Apparatus For Matching Multiple Images" filed on April 3, 2001, which is assigned to the same assignee as the present patent application and the teachings of which are herein incorporated by reference.

Background of the Invention

1. Field of the Invention

This invention generally relates to the field of image processing systems and methods, and more particularly relates to methods of recovering depth information associated with elements in a base image corresponding to multiple reference image views.

2. Description of Related Art

Image processing systems have attempted to process multiple images (views of a scene) to identify common image features across the different images (different views of a scene), such as to create three-dimensional (3-D) digital content by analyzing the multiple views of the scene. A main problem has been how to determine depth information for image feature elements, such as pixels, of a base image (view of a scene). When a scene is viewed from multiple cameras—one of them chosen as the base view, others as reference ones—the depth information of the scene for image features, such as pixels, in the base view's image plane can be recovered, based on the correspondence relationship between pixels in the base view and in the reference views.

The terminology "binocular stereo" refers to the case where two cameras are used, arranged in parallel to one another, as shown in FIG. 1. The distance between the two cameras is often called a baseline. The terminology "multi-baseline stereo" means the usage of multiple horizontally or vertically arranged cameras, also parallel to each other, as shown in FIG. 2.

Recent development in Image-Based Rendering and Modeling has raised the interest on multi-baseline stereo in the vision community. As a result of using multiple cameras arranged in multiple baselines two advantages are gained over binocular

stereo methods, i.e., a decrease of matching ambiguity, and an increase of reconstruction precision.

The basic problem of stereo, regardless how many cameras are used or how they are positioned, is to find the depth value of the 3 dimensional (3-D) scene point seen at each pixel of a base image, using other images as references. To accomplish this, for each pixel in the base image, its corresponding pixels (projections of the same scene point) in the reference images need to be identified. This correspondence problem can be very difficult to solve and impractical to implement, especially for more than a very small number of cameras, e.g., more than about 2 or 3 cameras. In the case of binocular or multi-baseline stereo, the task of establishing correspondence may be simplified in the sense that the corresponding pixels locate on the same horizontal or vertical scan line as the pixel in the base image. This makes possible representing the correspondence with a scalar, or *disparity*.

However, these binocular or multi-baseline stereo methods typically require special camera setup to achieve a common image plane, so that the cameras are necessarily coplanar and parallel. These methods unfortunately are also limited in the amount of coverage area of a particular scene. Further, these methods include restrictions on the camera placement that tend to complicate the overall image capture process, increase the cost of image capture, and generally make these methods impracticable for more complicated set ups, e.g., with more than a small number of

cameras. Typically, in these methods, either a mechanical device is used to ensure the cameras are collinear, or a mathematical process called rectification is performed to correct the mechanical misalignment. Lastly, the accuracy and reliability of results using these prior art methods would tend to be undesirable for serious commercial applications.

Therefore a need exists to overcome the problems with the prior art as discussed above, and particularly for a method and apparatus that can more successfully recover depth information for elements in a base image corresponding across multiple reference images.

Summary of the Invention

According to a preferred embodiment of the present invention, an image processing system comprises a memory; a controller/processor electrically coupled to the memory; an image matching module, electrically coupled to the controller/processor and to the memory, for providing a plurality of seed pixels that represent 3-D depth of the plurality of pixels in the base image view of a scene by matching correspondence to a plurality of pixels in a plurality of images representing a plurality of views of the scene; and a propagation module, electrically coupled to the controller/processor and to the memory, for tracing pixels in a virtual piecewise continuous depth surface by spatial propagation starting from the provided plurality of seed pixels in the base image by

using the matching and corresponding plurality of pixels in the plurality of images to create the virtual piecewise continuous depth surface viewed from the base image, each successfully traced pixel being associated with a depth in the scene viewed from the base image.

5

The image processing system preferably comprises at least one camera interface, electrically coupled to the controller/processor, for sending image information from at least one camera to the controller/processor.

10

According to a preferred embodiment of the present invention, the controller/processor, the memory, the image matching module, and the propagation module, are implemented in at least one of an integrated circuit, a circuit supporting substrate, and a scanner.

15

Brief Description of the Drawings

FIG. 1 is a block diagram illustrating a prior art binocular stereo arrangement of cameras.

20

FIG. 2 is a block diagram illustrating a prior art multi-baseline stereo arrangement of cameras.

FIG. 3 is a block diagram illustrating a multi-plane stereo arrangement of cameras, according to a preferred embodiment of the present invention.

FIG. 4 is a perspective view illustrating an exemplary geometrical relationship
5 between a set of image planes.

FIG. 5 is a diagram illustrating a potential relationship between images from two cameras.

10 FIG. 6 shows three corresponding images of a scene.

FIG. 7 shows a depth map of the center image of FIG. 6.

15 FIG. 8 shows two views of a reconstructed house model.

FIG. 9 shows four depth maps comparing a binocular stereo method vs. a Multi-Plane stereo method.

FIG. 10 is a block diagram illustrating an exemplary 3-D image processing
20 system according to a preferred embodiment of the present invention.

FIG. 11 is a more detailed view of the 3-D image processing system of FIG. 10.

FIG. 12 is an operational flow diagram illustrating an exemplary operational sequence for the 3-D image processing system shown in FIGs. 10 and 11, according to a preferred embodiment of the present invention.

5

Description Of The Preferred Embodiments

According to a preferred embodiment of the present invention, a new system and method utilizes a multi-plane stereo (MPS) arrangement of multiple cameras for recovering depth information from a scene. The multiple cameras advantageously are not limited to a coplanar and linear arrangement. See for example the arrangement shown in FIG. 3. According to a preferred embodiment of the present invention, there is no provision as how the cameras should be placed. Each image plane (associated with a particular camera) can be different than the other image planes (associated with the other cameras). Thus, a preferred method is called "multi-plane" stereo.

15 There are three main advantages with MPS. First, the camera arrangement covers a much larger area of a scene than the other two methods, i.e., the binocular stereo approach and the multi-baseline stereo approach. Second, reducing restrictions on the camera placement means simplification of the overall image information capture process. This makes a commercial application much more practicable and
20 commercially viable.

With the other two prior art methods, either a mechanical device is used to ensure the cameras are collinear, or a mathematical process called rectification is performed to correct the mechanical misalignment. Both of these constraints add tedious extra steps to an overall process of image capture. Additionally, since the multiple camera planes
5 extend in a three-dimensional region, more accurate and reliable results can be obtained.

In addition to binocular stereo and multi-baseline stereo, there has been some existing research on MPS that is mostly based on a regularization theory that suggested
10 solving a Partial Differential Equation (PDE) of the following form:

$$\frac{\partial z}{\partial t} = E_z(u, v, z, z_u, z_v)$$

where E is a function that encodes the similarity of a 3-D point's projection in the base image and its projections in the reference images. In other words, under this theory, z is thought of as a time-varying function governed by the above PDE. The final depth
15 corresponds to the equilibrium state of the evolution. This type of approach will tend to be more costly to implement and will be unable to effectively handle depth discontinuities.

A method, according to a preferred embodiment of the present invention, rather
20 than compute depth at individual pixels, defines a surface, called a *depth surface*, over an entire base image. This takes advantage of the cohesiveness of opaque object

surfaces. This method, in part, takes advantage of a *coherence principle*. What the principle implies is that most pixels of an image correspond to points of object surfaces where depth vary smoothly. Therefore, locally, depth recovery of neighboring pixels should facilitate each other.

5

Note that a discretized depth surface can be considered embedded in a volume. Each voxel can represent a possible match and is assigned a scalar value which encodes the quality of the match. The disparity surface comprises voxels with locally maximal values (i.e. best matching quality). The volume is thought of as a 6-connected
 10 network of voxels plus a multi-connected source and a sink. Each edge is associated with a scalar, a capacity, that is defined to be the matching cost. Since a likely match has a low matching cost, the corresponding edge capacity will be low and that edge is likely to be saturated by the maximum-flow. Inversely, a high matching cost yields a high capacity edge that is unlikely to be saturated. The set of edges that are saturated
 15 by the maximum-flow represent a minimum-cut of the graph. This cut effectively represents the depth surface.

To represent multi-plane stereo 3-D depth information as a surface for an entire base image, a suitable function E is defined over an open subset of the base image
 20 which minimize:

$$\iint E(u, v, z(u, v)) du dv \quad (1)$$

where E is the function that encodes the matching cost, and $z(u, v)$ is the depth surface sought. The corresponding Euler-Lagrange equation is readily obtained:

$$E_z(u, v, z(u, v)) = 0 \quad (2)$$

where E_z is the partial derivative of E with respect to z . A method that computes z separately for each pixel (u, v) can be deemed as a brute-force solution to (2), one of whose shortcomings is unable to model the surface cohesivity mentioned earlier.

A different method considered the function $z(u, v)$ also a function of time: $z(u, v, t)$, and solved the PDE:

$$\frac{\partial z}{\partial t} = E_z \quad (3)$$

with some initial condition: $z_0 = z(u, v, 0)$. One can imagine the process that a surface evolves over time, starting from the initial state specified by z_0 , governed by the PDE, shown by Equation (3) above, and converging to the true solution. One potential problem is that portions of the evolving surface could remain stuck in local minima due to image noise and repetitive intensity texture. To circumvent this problem, an image depth recovery algorithm can solve a regularized version of the problem:

$$\iint E(u, v, z, N) du dv \quad (4)$$

where $N = \frac{z_u \times z_v}{|z_u \times z_v|}$ is the normal of the surface. The surface normal is also used to specify a homography induced by the tangent plane which improves the correlation of

matching pixels. The assumption in the variational approach would be that $z(u, v)$ is smooth everywhere, which is often untrue, especially when there are multiple objects in the scene that occlude each other.

5 According to a preferred embodiment of the present invention, the Multi-Plane Stereo problem can be solved using a variational viewpoint. That is, the depth surface is determined by solving Equation (2) above. However, rather than cast it into a PDE problem where the depth surface evolves in time, a preferred method according to the present invention traces the surface using an approach the inventor calls spatial
10 propagation from known seed points in the object space. The advantage of this spatial propagation approach is that it does not require a depth surface to be smooth everywhere. In the present approach, partial differentiation on z is not explicitly performed which relieves the need on smoothness. Additionally, surface continuity is assumed as long as the propagation proceeds. The latter stops when discontinuity
15 happens.

A preferred method, according to the present invention, is based on the observation that a depth function can be calculated over a piecewise continuous surface over I_0 – a first image plane viewed from a base view. The preferred method
20 traces the surface calculated by a function resulting in a parameter g , the function starts from some predefined seed points on the surface.

**Generating Seed Points By: Single Image Feature Detection,
Two-View Matching Based On Cross-Correlation, Two-View Iterative
Refinement, And M-View Robust Matching.**

The seed points are, advantageously, obtained as the result of image matching
5 methods such as taught in U.S. Patent Application No. 09/825,266, entitled "Methods
And Apparatus For Matching Multiple Images", filed on April 3, 2001, which is
commonly owned and the entire teachings of which are hereby incorporated by
reference.

10 Since g does not depend on time, a simpler form of equation is solved for g at
each pixel position (u_j, v_j) :

$$E_g(u, v, g(u, v)) = 0. \quad (5)$$

15 For example, let $S=(u_j, v_j, g_j)$ be a seed point, i.e. the depth parameter at pixel (u_j, v_j) is g_j . S is then propagated to its four-neighbors at (u_{j-1}, v_j) , (u_{j+1}, v_j) , (u_j, v_{j-1}) and (u_j, v_{j+1}) . At each of these neighbors, the depth parameter is determined by solving Equation (5) using g_j as the initial value. Propagations from all seeds run concurrently. When two fronts meet, the one with lower cost prevails, the other one stops.
20 Propagation at a pixel also stops when the cost exceeds a certain threshold, or an image border is reached. Note that at a boundary of two surface pieces with different depths, the one with higher intensity texture tends to overshoot. Overshooting of a

surface propagation front means wrong correspondences and the matching cost increases sharply which causes it being surpassed by the propagation front of the other surface piece before expanding too far. Thus discontinuity is maintained. The more reference images used, the shorter overshooting, thus the better location of the depth discontinuity edge. Since the propagation actually happens in the object space, it is called spatial propagation.

According to a preferred method utilizing Multi-Plane Stereo, in general an image 3-D reconstruction system starts with $m+1$ images I_i ($i=0\dots m$) and the associated camera projection matrices P_i , and designates one of the images as the base image and the rest as reference images. The system then computes the depth for each pixel in the base image by means of correspondence in the reference images.

According to an exemplary preferred embodiment, in the following discussion with reference to FIG. 4, assume the first image 402 (with index $i=0$) is the base image 402. C_0 is defined as the base camera's center of projection (COP) 404. The image feature designated by (u, v) , in this example, is a pixel on the base image plane 406, and $L(u, v)$ 408 is the unit vector from C_0 to the point corresponding to (u, v) . The position of the 3-D point X 410 that projects to (u, v) can be represented as $C_0 + gL(u, v)$. The parameter g is related to the depth of X 410 by $z = g(L \cdot A_0)$ where A_0 is the unit vector representing the principal axis of the base camera. The parameter g is referred to herein as the depth parameter. Accordingly, a preferred method repeats the

calculation discussed above for other pixels in the base image 402 to compute the function $g(u, v)$ and obtain the parameter g for the other pixels.

A common definition for the matching cost E that appears in Equation (1) comes from the observation that when X 410 is back-projected to the reference images 412, 414, each back-projection $I_i(u', v')$, $i=1\dots m$, should reveal a similar intensity to that of $I_0(u, v)$, assuming perfect Lambertian surfaces and constant lighting. A simple way of describing the similarity of $I_i(u', v')$ and $I_0(u, v)$ is the squared difference $D_i = [I_i(u', v') - I_0(u, v)]^2 / 2$. Use the short hand $(u', v') = P_i(X)$ to denote that (u', v') is the projection of X in the i^{th} view. Then

$$D_i = [I_i(P_i(X)) - I_0(u, v)]^2 / 2 = [I_i(P_i(C_0 + gL(u, v))) - I_0(u, v)]^2 / 2.$$

Thus, one way of capturing the previous observation is to define E as the sum of D_i over all reference images I_i :

$$E = \sum_{i=1}^m [I_i(P_i(C_0 + gL(u, v))) - I_0(u, v)]^2 / 2. \quad (6)$$

It follows, then,

$$E_g = \sum_{i=1}^m [I_i(P_i(X(u, v, g))) - I_0(u, v)] \left(\nabla I_i \cdot \frac{\partial P_i}{\partial g} \right) = 0. \quad (7)$$

At this point, the solution to Equation (7) is postponed. The fact that Equation (7) is non-linear on g thus requires initial values almost immediately suggests a recursive

propagation scheme for a discrete solution to $g(u, v)$ over a regular grid. This discrete solution approach is a main focus for an implementation of a system according to a preferred embodiment of the present invention. Additionally, an improved similarity function for the discrete case will be discussed below.

5

Note a formulation, according to a preferred embodiment of the present invention, the correspondence problem is addressed implicitly with the help of back-projections of a 3-D point in question. This is in contrast with the traditional way the problem has been presented which does not involve 3-D points. Instead, they have been computed
10 in a subsequent step after the correspondence has been established.

Discretization

We are interested in a discrete solution of $g(u, v)$ over a regular grid (u_j, v_j) , $u_j=0, \dots, W-1$, $v_j=0, \dots, H-1$, where W and H are respectively horizontal and vertical number of
15 pixels of a base image. Given any u_j and v_j within their range, plugging them into Equation (7) results in a non-linear equation on g whose solution is the depth parameter at pixel (u_j, v_j) , denoted as g_j . Without an initial value for g , one has to resort to a global search on the ray $C_0+gL(u_j, v_j)$. Now, if it happens that one of (u_j, v_j) 's four-neighbors has a known depth parameter g' , then based on object surface cohesivity, it is likely that $g_j \approx$
20 g' . In other words, one could use g' as an initial value for g_j and refine it later. Applying this analysis recursively leads to the conjecture that as long as there exist some seed pixels whose depth parameters are known, the parameter surface can be traced.

When multiple object surfaces with depth discontinuities exist in the scene, theoretically, each piece of surface needs at least one seed.

Let $S=(u_j, v_j, g_j)$ be a seed point, i.e. the depth parameter at pixel (u_j, v_j) is g_j .

5 Recall from our discussion above that seed points are preferably obtained as the result of image matching methods such as taught in U.S. Patent Application No. 09/825,266, entitled "Methods And Apparatus For Matching Multiple Images", filed on April 3, 2001, which is commonly owned and the entire teachings of which are hereby incorporated by reference.

10

Then, S is propagated to its four-neighbors at (u_{j-1}, v_j) , (u_{j+1}, v_j) , (u_j, v_{j-1}) and (u_j, v_{j+1}) . At each of these neighbors, the depth parameter is determined by minimizing Equation (7). Propagations from all seeds operate concurrently. When two fronts meet, the one with lower cost prevails, the other one stops. Propagation at a pixel also
15 stops when the cost exceeds a certain threshold, or an image border is reached. The propagation occurring in the object space is called *spatial propagation*.

It is observed that, at a boundary of two surface pieces with different depths, the one with higher intensity texture tends to overshoot. Since overshooting means wrong
20 correspondences, its matching cost increases sharply which causes it being surpassed by the propagation front of the other surface piece before expanding too far. Thus

discontinuity is maintained. The more reference images used, the shorter overshooting, thus the better location of the depth discontinuity edge.

A Better Similarity Function In Discrete Case

5 The intuition behind Equation (6) is that intensity values at projections of X in all reference images should be close to $I_0(u, v)$. But the effect of discretization is not considered – we may not actually have a sample at $P_i(C_0 + gL(u_j, v_j))$ unless it happens to be a grid point in I_i . To cope with errors caused by discretization, it is more appropriate to consider square windows of certain size $\omega \times \omega$ that surround the putative
10 corresponding pixels, and define the matching cost as the normalized cross-correlation (NCC) between the windows:

$$E(g) = \sum_{i=1}^m \left\{ 1 - \text{NCC}_{\omega} [I_i(\bar{u}(g), \bar{v}(g)), I_0(u_j, v_j)] \right\} \quad (8)$$

where (\bar{u}, \bar{v}) is the grid point closest to $P_i(C_0 + gL(u_j, v_j))$, and the windows are respectively centered at (\bar{u}, \bar{v}) in I_i and (u_j, v_j) in I_0 . In other words, instead of finding exact fractional
15 pixels that correspond, which normally do not exist in a discrete case, windows with the maximal normalized cross-correlation are corresponded.

Rectification

When used as a similarity measurement for two images, cross-correlation can be
20 improved if the images are rectified – the corresponding epipolar lines become collinear

scanlines. Rectification is optionally utilized in a 3-D image reconstruction system according to a preferred embodiment of the present invention.

Rectification undergoes two steps: i) rotating both cameras around their center-
 5 of-projections so that their image planes are parallel to each other as well as parallel to the baseline; ii) adjusting the focal length of one or both cameras so that their image planes become coincident. After the rectification, both cameras have same focal length and orientation, and the common image plane is parallel to the baseline. Obviously, there is no unique solution here. Usually, one that minimizes image distortion is sought.

10 Let $K_i = \text{diag}(f_i, f_i, 1)$ and $R_i, i = 0 \dots m$, be the intrinsic and orientation matrices of the given cameras. Consider a pair of cameras consisting of the base camera and a reference one. A preferred method according to the present invention finds the rectification matrices for this pair. Let $K = \text{diag}(f, f, 1)$ and R be the intrinsic and
 15 orientation matrices of the rectified virtual cameras (no translation is involved since the cameras are only rotated). The rectification matrix is $M_i = KRR_i^{-1}K_i^{-1}$. Since the baseline between the base and the i^{th} cameras is known, what remain to be determined are the focal length f and the principal axis A of the virtual cameras. They are given below:

$$\begin{aligned} f &= (f_0 + f_i)/2 \\ A &= \text{normalize} \{ [B_{0i} \times (A_0 + A_i)] \times B_{0i} \} \end{aligned} \quad (9)$$

where B_{0i} is the vector representing the baseline. In Equation (9), A is undefined if $B_{0i} \times (A_0 + A_i) = 0$. In a typical MPS setup, it is expected that all cameras have similar principal axial directions, i.e. $A_0 \approx A_i$. Therefore the above singular condition implies $B_{0i} \approx A_0 \approx A_i$. This corresponds to a situation, as illustrated in FIG. 5, where one camera (thus the epipole) is in the view of the other one. In this situation, however, the depth estimation close to the epipole is highly unreliable. In practice, such kind of degenerate configuration should be avoided. When it happens (within some threshold), rectification should not be performed.

Some Remarks About A Preferred Implementation

1) By explicitly involving the camera projection matrices in the similarity measurement, maximal camera placement flexibility is achieved.

2) Compared to the often-used SSD (Sum of Squared Difference) function, NCC normalizes image intensities.

3) When g varies, $C_0 + gL(u_j, v_j)$ spans a ray emitting from C_0 and passing the grid point (u_j, v_j) . The projection of the ray in the i^{th} reference image – $P_i(C_0 + gL(u_j, v_j))$ – is an epipolar line in that image. Therefore, the epipolar constraint is implicitly enforced in Equation (8) between the base image and each reference one. This is a very strong condition, especially when the cameras are generally placed.

4) In Equation (8), when $m=1$, MPS reduces to binocular stereo. When all the cameras are horizontally aligned, i.e. either $\bar{u} = u_j$ or $\bar{v} = v_j$ for (\bar{u}, \bar{v}) in all the

reference images, MPS reduces to multi-baseline stereo. In this sense, the proposed formulation, according to the present invention, is the most general one because others can be treated as special cases.

5

Discussion Of Some Experimental Results

In this section, some experimental results are presented. The first example consists of five images, three of them shown in FIG. 6. Among them, the center one is the designated base image. A depth map of the center image of FIG. 6 is shown in FIG. 7 where darker pixels correspond to shorter depth, and where white pixels, on the other hand, represent areas where depth cannot be estimated. Two such areas are observed. First are the background and the roof of the house where intensity texture is insufficient. Second is the frontal part of the carpet where, due to foreshortening, NCC is low (therefore cost is high).

10

15

FIG. 8 shows two views of the reconstructed house model. Observe that discontinuities at both sides of the house walls are preserved. Also observe the shapes on the sidewall which indicate doors and windows.

20

In FIG. 9, depth maps are shown as shaded 3-D surfaces where the XY-plane represents the image coordinate system, and the Z-axis represents depth. FIG 9 shows four views, FIG. 9A and FIG. 9B represent depth surface views using a binocular

stereo method, while FIG. 9C and FIG. 9D represent depth surface views using an MPS method. Visually, it is seen that the depth surface from two-view binocular stereo is much bumpier than that from MPS. Particularly, notice the sidewall of the house and the portion of the carpet close to the right image border, as indicated by the arrows 902, 904.

Some Potential Benefits In Commercial Applications

Once the depth is obtained, two-dimensional images can be enhanced with three-dimensional information. This, for example, allows augmenting images in several ways. As a first example, by combining the original intensity image with its depth image, a new kind of image is generated. In this depth-enhanced image, as a second example, a synthetic object can be inserted into the image which forms different mutual occlusion with those existing objects in the original scene. Other applications include augmenting images by adding synthetic lights or changing the geometric shapes of the objects in the images. For instance, before conducting a plastic surgery to a person, his or her photos are taken and 3-D geometry is formed from the depth information. Then, the photos can be modified showing the patient's new nose or new breasts until the patient's expectations are satisfied. The collected information is then forwarded to the surgeon to help the surgeon more accurately conduct the surgical procedure.

Exemplary System Implementation

According to a preferred embodiment of the present invention, as shown in FIG. 10, an exemplary 3-D image processing system 1000 comprises a set of digital (still or video) cameras 1002, three cameras 1004, 1006, 1008, being shown, which are arranged with different poses and are electronically synchronized such as via an electrical signal bus 1010. At any time instant, the cameras 1002 generate a set of images 1012, such as three images 1014, 1016, 1018, being shown for the three respective digital capture interfaces 1015, 1017, 1019, for the three cameras 1004, 1006, 1008. Each of the set of images 1012 deviates from the other images in the set of images 1012 by camera relative motion. For example, the first image 1014 and the second image 1016 can deviate from one another by a distance between corresponding feature points found on both images 1014, 1016, due to the different poses of the cameras 1004, 1006, relative to a scene 1026. This camera relative motion between the two images 1014, 1016, can be represented by a motion vector between feature points that correspond (i.e., that match) between the two images 1014, 1016. Additionally, although still cameras 1004, 1006, 1008, and a still scene 1026 are shown in this example, it should become obvious to one of ordinary skill in the art in view of the teachings herein that any combination of still and/or moving scene 1026 and/or cameras 1004, 1006, 1008, can be represented in accordance with alternative embodiments of the present invention. For example, a moving object scene 1026 and utilizing still cameras 1004, 1006, 1008, may be perfectly desirable for certain applications of the present invention. Therefore, the term camera relative motion, as

used herein, is intended to broadly cover all such alternative embodiments of the present invention wherein any combination of still and/or moving scene 1026 and/or cameras 1004, 1006, 1008, can be represented.

5 The three respective digital capture interfaces 1015, 1017, 1019, are communicatively coupled to a computer system (not shown in FIG. 10). The set of images 1012 is then processed by the hardware 1020, the computer system (not shown), and the software 1022 of the system 1000 to output 3-D image information 1024 of the scene 1026 observed by the set of cameras 1002. The software 1022
10 preferably comprises a point detection and matching module 1028, as taught in U.S. Patent Application No. 09/825,266, entitled "Methods And Apparatus For Matching Multiple Images", filed on April 3, 2001, which is commonly owned and the entire teachings of which are hereby incorporated by reference. This image matching module 1028 provides the seed points for further processing to extract the 3-D depth
15 information. A 3-D image reconstruction module 1030, as will be discussed in more detail below, provides additional processing of the image information after the image feature points have been detected and matched across views to capture a 3-D image of the scene 1024 that, for example, can be displayed (such as via a display) to a user of the system 1000.

20

FIG. 11 illustrates a more detailed view of the 3-D image processing system 1000 of FIG. 10, according to a preferred embodiment of the present invention. Each

of the digital capture interfaces 1015, 1017, 1019, includes respective image capture memory 1104, 1106, 1108, for storing a captured image 1014, 1016, 1018. The digital capture interfaces 1015, 1017, 1019, are communicatively coupled to an input/output interface 1110 of a 3-D image reconstruction computer system 1102. Additionally, the electrical signal bus 1010 is communicatively coupled to the input/output interface 1110. The 3-D image reconstruction computer system 1102 comprises a controller/processor 1112 that is electrically coupled to data memory 1114 and to program memory 1116. The controller/processor 1112 is also electrically coupled to a user interface 1118 that presents information to a user, such as via a monitor display (not shown), and receives user input from the user such as via a keyboard (not shown) and a mouse (not shown).

The data memory 1114 includes an image memory 1120 for storing image information. The image memory 1120 comprises data structures for a seed list 1150, a Propagate memory 1152 to keep track of front propagation results for each piece of depth surface being propagated, a Window memory 1154 to keep track of windows surrounding pixels for propagating each piece of depth surface, and a 3-D Surface memory 1156 that keeps track of the propagated pieces of depth surface. These data structures are used by the 3-D image reconstruction functional module 1028 as will be discussed in more detail below. Additionally, the data memory 1114 includes a parameter memory 1122 where the 3-D image reconstruction computer system 1102 stores configuration parameters for the 3-D image processing system 1000.

5 The program memory 1116 provides computer program instructions for the controller/processor 1112 for performing operational sequences for the 3-D image processing system 1000, according to the preferred embodiments of the present invention. The program memory 1116 includes four functional modules. The four functional modules are as follows: an image matching module 1028, a propagation module 1132 for computing the propagation of depth parameter for surface pieces, a cost compare handler 1134 for computing the relative cost of two surface pieces being propagated and meeting at a boundary, and a rectification handler 1136 that can be optionally used to operate on rectification matrices of images to improve the image processing by cross-correlation. These four functional modules will be discussed in more detail below.

10 Additionally, the 3-D image reconstruction computer system 1102 preferably includes a drive 1140 for receiving the computer readable medium 1142. This provides a means of transferring of information with the 3-D image reconstruction computer system 1102. For example, computer programs (and updates thereto) can be provided to the 3-D image reconstruction computer system 1102 and stored in the program memory 1116 via the computer readable medium 1142 in a manner well known to those of ordinary skill in the art. Additionally, image information and related parameters can be transferred between the computer readable medium 1142 and the data memory 1114.

According to a preferred embodiment of the present invention, the point detection and image matching module 1028 operates in the 3-D image reconstruction computer system 1102 and is stored in the program memory 1116. The image matching module 1028, according to one embodiment of the present invention, operates on image information in a series of four operational stages that progressively improve the correspondence of image information across a plurality of images representative of a scene. As a result of the fourth stage, i.e., a multiple-view robust matching handler, the correspondence of the image information across the plurality of images (views) is significantly improved over known prior art systems. The resulting image information that corresponds across the plurality of views provides image features (such as pixels) as seeds for the 3-D image reconstruction computer system 1102 to extract the 3-D image depth information as discussed above, according to the present invention.

According to a preferred embodiment of the present invention, significant portions of the 3-D image processing system 1000 may be implemented in integrated circuits. For example, functional components of the 3-D image reconstruction computer system 1102 may be implemented in at least one integrated circuit. Similarly, significant portions of the digital capture modules 1015, 1017, 1019, can be implemented in at least one integrated circuit.

According to alternative embodiments of the present invention, the 3-D image processing system 1000 may be implemented, for example, in devices such as three-dimensional scanners, facsimile machines, video communication equipment, and video processing equipment.

5

Referring to FIG. 12, the controller/processor 1112 enters the operational sequence at step 1202, and initializes system parameters in the parameter memory 1122. The controller/processor 1112 then operates according to the image matching module 1028 to provide seeds, at step 1208. The image matching module 1028, as
10 discussed above, utilizes a series of functional modules to match image features across the plurality of images from the cameras 1004, 1006, 1008, resulting in a set of image features, such as pixels, that are stored in the seed list memory 1150, at step 1208.

The operational sequence then repeats, at steps 1210, 1211, 1212, 1213, 1214,
15 1215, and 1216, to propagate depth surface fronts from the seed points until the entire depth surface computations for the base view has been completed. At step 1212, if the cost of a newly propagated pixel is too high, the controller/processor 1112 abandons it and directly goes to further propagation. At step 1213 when a boundary of two surface pieces is reached, the matching costs of the piece surfaces are compared, at step
20 1214, and the propagation of the higher cost surface is stopped. After the entire depth surface has been computed, at step 1216, the operational sequence then exits, at step 1218.

3-D Image Reconstruction System Realization

5 The present invention can be realized in hardware, software, or a combination of hardware and software. A system according to a preferred embodiment of the present invention can be realized in a centralized fashion in one computer system, or in a distributed fashion where different elements are spread across several interconnected computer systems. Any kind of computer system - or other apparatus or integrated circuit adapted for carrying out the methods described herein - is suited. A typical combination of hardware and software could be a general purpose computer system
10 with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein.

15 The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which - when loaded in a computer system - is able to carry out these methods. Computer program means or computer program in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another
20 language, code or, notation; and b) reproduction in a different material form.

Each computer system may include, inter alia, one or more computers and at least a computer readable medium allowing a computer to read data, instructions, messages or message packets, and other computer readable information from the computer readable medium. The computer readable medium may include non-volatile
5 memory, such as ROM, Flash memory, Disk drive memory, CD-ROM, and other permanent storage. Additionally, a computer medium may include, for example, volatile storage such as RAM, buffers, cache memory, and network circuits. Furthermore, the computer readable medium may comprise computer readable information in a transitory state medium such as a network link and/or a network interface, including a wired
10 network or a wireless network, that allow a computer to read such computer readable information.

The 3-D image processing system 1000 according to the present invention provides significant advantages over the known prior art. The present system 1000 can
15 be much more accurate and efficient at reconstructing 3-D image information for a base view relative to multiple reference views. The 3-D image processing system 1000 according to the present invention provides a practicable approach to capturing 3-D image information from across multiple images of a scene.

20 Accordingly, due to the remarkable efficiency of the embodiments of the present invention, an implementation in an integrated circuit (IC) chip is very feasible and desirable. Generally, a circuit supporting substrate and associated circuits, such as

provided by an IC, a circuit board, and a printed circuit card, and other similar
embodiments, and including the functional modules according to the present invention
as discussed above, can provide a modular solution for enabling a computer system to
benefit from the very accurate 3-D image processing methods according to the present
5 invention. Such electronic devices as a three dimensional scanner and a 3-dimensional
video image capture system are commercially feasible. Additionally, since the system
according to the present invention can beneficially utilize many more cameras, e.g.,
more than two or three cameras, then the 3-D image capture methods of the present
invention can advantageously operate both with a still object in a scene as well as with
10 a moving object in a scene.

Although specific embodiments of the invention have been disclosed, those
having ordinary skill in the art will understand that changes can be made to the specific
embodiments without departing from the spirit and scope of the invention. The scope
15 of the invention is not to be restricted, therefore, to the specific embodiments, and it is
intended that the appended claims cover any and all such applications, modifications,
and embodiments within the scope of the present invention.

What is claimed is:

20